# The Peter Principle: A Theory of Decline

Edward P. Lazear

*Hoover Institution and Stanford University*

Some have observed that individuals perform worse after being promoted. The Peter principle, which states that people are promoted to their level of incompetence, suggests that something is fundamentally misaligned in the promotion process. This view is unnecessary and inconsistent with the data. Below, it is argued that ability appears lower after promotion purely as a statistical matter. Being promoted is evidence that a standard has been met. Regression to the mean implies that future ability will be lower, on average. Firms optimally account for the regression bias in making promotion decisions, but the effect is never eliminated. Rather than evidence of a mistake, the Peter principle is a necessary consequence of any promotion rule. Furthermore, firms that take it into account appropriately adopt an optimal strategy. Usually, firms inflate the promotion criterion to offset the Peter principle effect, and the more important the transitory component is relative to total variation in ability, the larger the amount that the standard is inflated. The same logic applies to other situations. For example, it explains why movie sequels are worse than the original film on which they are based and why second visits to restaurants are less rewarding than the first.

The Peter principle states that workers are promoted to their level of incompetence (Peter and Hull 1969). One interpretation is that firms

Sherwin Rosen was my most important teacher, my valued colleague, and dear friend. Sherwin served on my thesis committee and taught me much of what I know. Throughout the 30 years that we were friends, Sherwin was a constant source of inspiration, wisdom, and kindness. A deep thinker who opened up a number of areas of research, Sherwin was interested in hierarchies and promotion, so this paper is very much in keeping with his research agenda and derives from it. This research was supported in part by the National Science Foundation. Useful comments were provided by many. Po-Han Fong, Thomas P. Gherig, Kevin J. Murphy, Torsten Persson, Paul Pfleiderer, Kathryn Shaw, Eskil Wadensjo, and Michael Waldman were particularly helpful.

[*Journal of Political Economy,* 2004, vol. 112, no. 1, pt. 2]

systematically make mistakes in their promotion decisions. Another, favored here, is that the decline in ability that is seen after promotion is the natural outcome of a statistical process that displays regression to the mean. Workers are promoted on the basis of having met some standard. To the extent that ability is the sum of both permanent and transitory components, those who meet the standard have expected transitory components that are positive. The expectation of the post-promotion transitory component is zero, implying a reduction in expected ability. Firms that understand the statistical process take this phenomenon into account by adjusting the promotion standard, but the result remains: Expected ability for those promoted is lower after promotion than before.[1]

There is substantial evidence of the Peter principle.[2] In addition to papers from the marketing and organizational behavior literature (see, e.g., Anderson, Dubinsky, and Mehta 1999), there are a number of findings in empirical labor economics that support the claim. In an early paper that used subjective performance evaluation, Medoff and Abraham (1980) reported that workers' subjective evaluation scores fell the longer they were on the job. In Lazear (1992), it was found that the coefficient of job tenure in a wage regression was actually negative. The longer a worker was in a particular job, given his tenure in the firm, the lower his wage. The reason presumably is that the better workers are promoted out of the job, so those with a given number of years of firm experience who have fewer years in a job are less likely to have gotten stuck in that job. Baker, Gibbs, and Holmstrom (1994) replicate this finding in their data, and Gibbs and Hendricks (2001) find that raises and bonuses fall with tenure.

The tone of the literature outside of economics is that there is something wrong with promotion dynamics, and this anomaly shows up as the Peter principle. (Indeed, the book written by Peter and Hull is entitled *The Peter Principle: Why Things Always Go Wrong*.) The approach taken here is different. The Peter principle results from optimal adjustment to decision making under uncertainty. It is argued that even when firms use exactly the right promotion rule, the Peter principle

[1] It is also true that those who are denied promotion do better after they are turned down than they did before the decision was made, for the same reason.

[2] Two recent papers (Faria 2000; Fairburn and Malcolmson 2001) on this topic use an approach very different from this paper and from each other. The Peter principle is a by-product of using promotion to solve a moral hazard problem in Fairburn and Malcolmson. Rather than motivate through money, which induces influence activity, firms choose promotion because then managers must live with the consequences of their decisions. Too many workers are promoted under certain circumstances, resulting in a Peter principle effect. In Faria's paper, workers have two skills. Those who are good at one are necessarily less good at another when on the frontier. Faria argues that this is what is meant by the Peter principle.

effect will be observed. The fact that reversals of promotion decisions are rare is more compatible with the view that they were correct in the first place than that something went wrong.

More often, the Peter principle is interpreted in a multifactor context. Individuals who are good in one job are not necessarily good in the job into which they are promoted. As a result, individuals appear incompetent in the job in which they settle. To obtain this result, it is merely necessary to make a slight modification in the regression to the mean structure. Here, general ability is combined with a job-specific ability to produce output. Regression to the mean results because positive readings on the job-specific component prior to promotion are uncorrelated with the job-specific component after promotion.[3]

The fact that promoted individuals are less able than their apparent prepromotion ability induces firms to adjust in two respects. First, firms select their promotion rule with the understanding that the prepromotion ability is a biased estimate of true ability for those who exceed some standard. Second, as the variance in the transitory component of ability rises relative to the variance in the permanent component, the adjustment factor becomes greater. In the typical case, the standard that one must exceed to obtain a promotion increases with the relative importance of the transitory component because the regression effect is larger when the variance of the transitory component is large.

The model presented below yields the following results:

1. Promoted individuals' performance falls, on average, relative to their prepromotion performance.
2. Firms that take the decline into account adjust their promotion rule accordingly, but this does not negate the observation that ability declines after promotion.
3. The importance of the Peter principle depends on the amount of variation in the transitory component relative to the permanent. The Peter principle is most pronounced when the transitory component is large.
4. The length of the prepromotion period depends on the ratio of transitory variation to permanent variation. As the transitory component becomes more important, firms lengthen the prepromotion period.
5. Movie sequels are systematically worse than the original on which they are based.
6. Follow-up visits to good restaurants provide poorer meals than the first sampling.

[3] The structure is a variant of the Jovanovic (1979*a*, 1979*b*) model that was modified and used in a context closer to this structure in Lazear (1986).

7. In the absence of learning effects, second-term elected officials are less effective than they were during the first term.

## I. Model

Let there be two periods. Each worker has a time-invariant component of ability, denoted $A \sim f(A)$, and a time-varying component of ability, denoted $\epsilon_1$ for period 1 and $\epsilon_2$ for period 2. Let the time-varying components be independently and identically distributed with density $g(\epsilon)$. The firm can observe $A + \epsilon_t$ in each period but cannot disentangle the time-varying component of ability from the permanent component. There are a variety of interpretations that are consistent with this specification. One can think of the $\epsilon_t$ as being a true transitory aspect of ability or just measurement error. Later, the interpretation of different jobs will be considered.

There are two jobs (two are sufficient), which we denote difficult and easy. An individual's productivity in the easy job is given by

$$\alpha + \beta(A + \epsilon_t)$$

and in the difficult job is given by

$$\gamma + \delta(A + \epsilon_t),$$

where $\alpha > \gamma$ and $\delta > \beta$. Thus it pays to assign a worker to the difficult job if and only if

$$A + \epsilon_t > x,$$

where $x \equiv (\alpha - \gamma)/(\delta - \beta)$. The situation and the crossing point that correspond to $x$ are shown in figure 1.[4] The setup seeks to capture the idea that the most able have a comparative advantage in the difficult job.

Assume that individual ability $f(A)$ is such that, in the absence of information, it pays to assign everyone to the easy job in period 1.[5] Intuitively, this assumption amounts to saying that most people are not well suited to the difficult job and that, in the absence of countervailing information, individuals are assigned to the easy job. With one excep-

[4] This production structure is similar to that used in a comprehensive analysis by Gibbons and Waldman (1999*b*), who also allow for transitory and permanent components with regression. The focus of their paper is earnings and promotion. Neither optimal decision making by firms given the transitory component nor strategic effort in response to promotion rules is central to their discussion.

[5] This amounts to assuming that

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}[\alpha + \beta(A + \epsilon)]\,dGdF > \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}[\gamma + \delta(A + \epsilon)]\,dGdF.$$

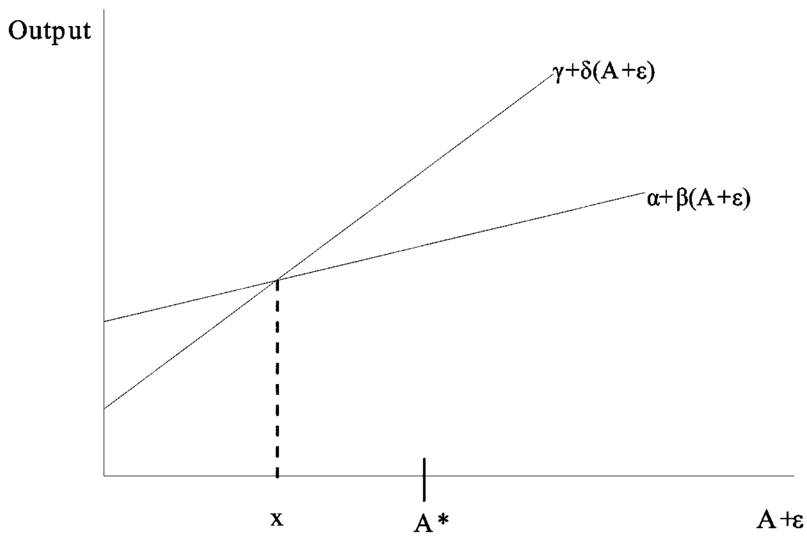


FIG. 1

tion, noted in the Appendix, it is sufficient to assume symmetric ignorance, so that workers are no better informed about their abilities than firms.

After the first period, firms obtain an estimate of $A$, namely $\hat{A} = A + \epsilon_1$. Since $\epsilon_1$ is the period 1 transitory component (either measurement error or transitory ability), it is $A$ and not $\hat{A}$ on which a promotion decision should be made. But $A$ is not observed, so firms are forced to base their decision on $\hat{A}$.

### A. *Workers Perform Worse after Being Promoted*

Firms must select some criterion level, $A^*$, such that if $\hat{A} > A^*$, the worker is promoted to the difficult job. If $\hat{A} < A^*$, the worker remains in his current job. It is now shown that workers who are promoted have levels of ability in period 1 that are higher, on average, than their ability in period 2.

First, note that the expectation of $\epsilon_1$ given that an individual is promoted is

$$\begin{aligned} \mathrm{E}(\epsilon_1 | A + \epsilon_1 > A^*) &= \int_{-\infty}^{\infty}\int_{A^*-A}^{\infty} \frac{1}{1 - G(A^* - A)} \epsilon g(\epsilon) f(A) d\epsilon dA \\ &= \int_{-\infty}^{\infty} \mathrm{E}(\epsilon | \epsilon > A^* - A) f(A) dA, \end{aligned}$$

which is positive since $f(A)$ is positive and the conditional expectation of $\epsilon$ given $\epsilon$ greater than any number is positive (because the unconditional expectation of $\epsilon$ is zero). Thus the conditional expectation of $\epsilon_1$ is positive among those who are promoted.

Now, in period 2, the expectation of the transitory component is

$$E(\epsilon_2 | A + \epsilon_1 > A^*) = 0$$

because $\epsilon_2$ is independent of $A$ and of $\epsilon_1$. As a result, for any promoted individual with ability $A$,

$$A + E(\epsilon_1 | A + \epsilon_1 > A^*) > A + E(\epsilon_2 | A + \epsilon_1 > A^*).$$

Thus expected ability falls for promoted individuals from period 1 to period 2.

Individuals who are promoted are promoted in part because they are likely to have high permanent ability,[6] but also because the transitory component of their ability is high. One of the reasons that academics tend to write better papers before they receive tenure is that they would not have received tenure had they not written the better-than-average papers. The point is obvious but is made graphic by the following example. Suppose that a firm promotes all individuals who can obtain three heads on three consecutive coin tosses. Only one in eight will be promoted. But when the firm asks the promoted individuals to repeat the feat, only one in eight will measure up. Seven out of eight will do worse than they did before being promoted. The reason is that all of the "performance" on the coin toss is transitory since tosses are independent.

As a general matter, the larger the transitory component is relative to the permanent component, the more important the Peter principle effect. If there were no transitory component, there would be no regression to the mean. Thus the importance of "luck" is positively associated with the force of the Peter principle.

### *B. The Promotion Rule*

Firms know in advance that there will be some expected fall in productivity among those promoted and adjust their promotion standard accordingly. Below, the general optimization problem for the firm is presented. Then, the way in which the rule operates is demonstrated by an example.

The firm's problem is to maximize profits (or worker utility) by selecting the job for each candidate with the highest expected value. Recall

[6] The notation $\epsilon_1$ and $A$ could be swapped in the discussion above to show that $E(A | A + \epsilon_1 > A^*) > E(A)$.

that individuals who have period 2 ability greater than $x$, defined above, would be assigned to job 2 were second-period ability known. The firm does not see $A$, but only $\hat{A}$, and must choose some criterion, $A^*$, such that it promotes workers whose observed ability in period 1 is greater than $A^*$. This is equivalent to promoting individuals when $A > A^* - \epsilon_1$. Thus the firm wants to choose $A^*$ so as to maximize

$$\max_{A^*} \int_{-\infty}^{\infty}\int_{A^*-\epsilon_1}^{\infty}\int_{-\infty}^{\infty} [\gamma + \delta(A + \epsilon_2)] f(A) g(\epsilon_1) g(\epsilon_2) d\epsilon_2 dA d\epsilon_1$$

$$+ \int_{-\infty}^{\infty}\int_{-\infty}^{A^*-\epsilon_1}\int_{-\infty}^{\infty} [\alpha + \beta(A + \epsilon_2)] f(A) g(\epsilon_1) g(\epsilon_2) d\epsilon_2 dA d\epsilon_1. \quad (1)$$

Because the expectation of $\epsilon_2$ is zero, (1) can be written as

$$\max_{A^*} \int_{-\infty}^{\infty}\int_{A^*-\epsilon_1}^{\infty} (\gamma + \delta A) dF dG + \int_{-\infty}^{\infty}\int_{-\infty}^{A^*-\epsilon_1} (\alpha + \beta A) dF dG. \quad (2)$$

The choice of $A^*$ depends on the distribution. However, two examples reveal that $A^*$ does not equal $x$ as a general rule. Instead, in typical cases, firms adjust $A^*$ upward. Knowing that worker ability in period 2 will differ from worker ability in period 1, firms usually set the bar higher than they would, where ability observed in period 1 carried over directly to period 2.

Actual solutions that provide intuition are available for given distributions. Consider, for example, the case in which $A$, $\epsilon_1$, and $\epsilon_2$ are all distributed normally, with mean zero and variance equal to one. Let $\alpha = 1$, $\beta = 0.5$, $\gamma = 0$, and $\delta = 1$. Then $x$, the ability level at which jobs produce equal value, is two since

$$\alpha + \beta(A + \epsilon) = \gamma + \delta(A + \epsilon)$$

for $A + \epsilon = 2$. However, $A^* = 4.01$. The firm sets its promotion standard more than two standard deviations higher than the crossing point in figure 1 because it understands that the worker's ability in period 2 is likely to be lower than it was in period 1 for the promoted group. As a result, the firm insists on a very high level of observed ability in period 1 in order to warrant promotion. Statements such as "tenure requires that the faculty member be the best in his or her field, having produced outstanding research" is a manifestation of the upward adjustment.

Consider the same example with a twist. Let the distribution of $A$ remain the same, namely, normal with mean zero and a standard deviation of one, but let the standard deviation of $\epsilon_1$ fall to 0.1. Then $A^*$ drops from 4.01 to 2.08. Although the firm still adjusts its promotion criterion upward from $x$, the adjustment is much smaller because the importance of the transitory component has been diminished. There is regression to the mean, but the regression that takes place is small relative to the amount in the prior example. When the standard deviation of $\epsilon$ is zero, the promotion standard is two, which is exactly $x$ as expected. Then the distribution of $\epsilon_1$ is degenerate, so that all observed ability in period 1 is permanent ability. The problem in (2) becomes

$$\max_{A^*} \int_{A^*}^{\infty} (\gamma + \delta A) f(A) dA + \int_{-\infty}^{A^*} (\alpha + \beta A) f(A) dA,$$

which has first-order condition

$$\frac{\partial}{\partial A^*} = -(\gamma + \delta A^*) f(A^*) + (\alpha + \beta A^*) f(A^*) = 0.$$

The solution is

$$\gamma + \delta A^* = \alpha + \beta A^*,$$

which is the crossing point, that is, $x$, in figure 1. When there is no transitory component, the firm simply promotes those whose permanent ability places them better in the difficult job than in the easy job.

It is possible to derive the relation between $A^*$ and $x$ in more general terms.[7] The first-order condition to (2) is

$$(\gamma + \delta A^* - \alpha - \beta A^*) \int_{-\infty}^{\infty} f(A^* - \epsilon_1) g(\epsilon_1) d\epsilon_1 =$$

$$(\delta - \beta) \int_{-\infty}^{\infty} \epsilon_1 f(A^* - \epsilon_1) g(\epsilon_1) d\epsilon_1.$$

The integral on the left-hand side is always positive, so whether $A^*$ exceeds $x$ or not depends on the sign of the integral on the right-hand side. Assume that both $f(\cdot)$ and $g(\cdot)$ are symmetric densities, and let $g(\cdot)$ be symmetric around zero and $f(\cdot)$ be symmetric around $\bar{A}$. Write the integral on the right side as

$$\int_{-\infty}^{0} \epsilon_1 f(A^* - \epsilon_1) g(\epsilon_1) d\epsilon_1 + \int_{0}^{\infty} \epsilon_1 f(A^* - \epsilon_1) g(\epsilon_1) d\epsilon_1.$$

Use a change of variable in the first integral of $u = -\epsilon_1$, and in the second, allow $u = \epsilon_1$. Because of symmetry, $g(u) = g(-u)$, so one can

[7] I am indebted to Wing Suen for this derivation.

write the two integrals as one:

$$\int_0^{\infty} u[f(A^* - u) - f(A^* + u)]g(u)du.$$

Suppose that the firm wants to promote fewer than 50 percent of the people, that is, that $x > \bar{A}$. (Recall that $x$ is the value such that $\alpha + \beta x = \gamma + \delta x$. If $x = \bar{A}$, then because of symmetry of the density functions, half of the population would have $A > x$ and half would have $A < x$.) Under such circumstances, $f(x - u) > f(x + u)$ (because $f$ is unimodal around $\bar{A}$). Thus the right-hand side of the first-order condition is positive for $A^* = x$, which implies that

$$\alpha + \beta A^* > \gamma + \delta A^*$$

or that $A^* > x$. Thus the firm adjusts the cutoff level upward when fewer than 50 percent of the workers are better suited to the difficult job than to the easy job.

The same reasoning applies in reverse. If more than 50 percent are to be promoted, then $x < \bar{A}$, which means that $f(x - u) < f(x + u)$. As a result, the right-hand side of the first-order condition is negative at $x$, which means that $A^*$ must be less than $x$ to satisfy optimality. Thus, when more than 50 percent are better suited to the difficult job than to the easy job, the firm reduces the promotion cutoff below the one that would be optimal were there no error in period 1. (Actually, under such circumstances, workers would initially be assigned to the difficult job, and the standard would be one such that workers who fell below it would be demoted after the first period.)

The intuition is this. Although there is always regression to the mean, adjusting the promotion level upward reduces the probability that the firm will make a bad promotion decision. However, at the same time, the adjustment increases the probability that it will fail to promote a qualified worker; that is, it reduces the false positive while increasing the false negative error. Conversely, lowering the promotion cutoff reduces the probability that someone who erroneously was observed to be a poor worker is not promoted but increases the probability that the firm promotes too many bad workers. Thus there is a trade-off. When fewer than 50 percent are to be promoted, the expected cost of making a false positive error exceeds that of making a false negative error so that the criterion must be adjusted upward. To the extent that most hierarchies are narrower at the top than at the bottom, $A^* > x$ is probably the more typical case. Standards are adjusted upward.

*C. Another Interpretation of the Peter Principle*

Rosen (1986) presents a model of sequential promotions in which individuals are sorted by ability at each stage such that members of the entering class at each round of a tournament have equal (ex ante) ability. Rosen uses the model to determine the optimal compensation at each level to motivate workers. Sorting is also an issue because the pool of workers at successive rounds have higher ability than those at earlier rounds. The Rosen model is in some ways more general than the one described here because it allows for effort as well as ability differences. The focus is quite different, however, because neither the optimal promotion rule nor the worker's output over time is an important part of the analysis. It is likely that many of the Peter principle results that come out of this paper could have been derived in that important paper on sequential promotions.[8]

Still, the Rosen model does not fit one of the most common interpretations of the Peter principle, which is that workers are promoted to their level of incompetence because a worker who is good in one job is not necessarily good in a job one level up. Fine professors do not necessarily make good deans (although not all would interpret moving to the dean's job as a promotion).[9] A slight modification of the definitions above and some of the formulas permits this interpretation.

To see this, allow $\epsilon_1$ to be defined as the job-specific component of ability associated with the easy job and $\epsilon_2$ as the job-specific component of ability associated with the difficult job. Individuals are assigned to the easy job in period 1 for the reason given before: Most are better at the easy job, and in the absence of information, the easy job is the right assignment. After evaluation, $\hat{A}$ is observed and the worker is promoted or not. If he is not promoted, then his ability after promotion is $A + \epsilon_1$. If he is promoted, then his ability after promotion is $A + \epsilon_2$.

Under this interpretation, workers who are not promoted have output that remains constant over time and equal to $\alpha + \beta(A + \epsilon_1)$. Those who are promoted have output equal to $\gamma + \delta(A + \epsilon_2)$. The argument of subsection *A* holds:

$$\mathrm{E}(\epsilon_1|A + \epsilon_1 > A^*) = \int_{-\infty}^{\infty}\int_{A^*-A}^{\infty}\frac{1}{1 - G(A^* - A)}\epsilon g(\epsilon)f(A)d\epsilon dA$$

$$= \int_{-\infty}^{\infty}\mathrm{E}(\epsilon|\epsilon > A^* - A)f(A)dA,$$

[8] Gibbons and Waldman (1999*a*) and Prendergast (1999) survey the large literature on careers and promotions.

[9] Anderson et al. (1999) claim that the data reveal a Peter principle for sales managers because the skills needed by salespeople are generally distinctly different from those needed by sales managers.

which is positive since $f(A)$ is positive and the conditional expectation of $\epsilon$ given $\epsilon$ greater than any number is positive (because the unconditional expectation of $\epsilon$ is zero). But the expectation of $\epsilon_2$ is zero for promoted workers because $\epsilon_1$ and $\epsilon_2$ are uncorrelated. As a result, expected ability is higher before promotion than for promoted workers after promotion.

This does not necessarily imply that output is lower after promotion because workers are in different jobs. On the contrary, if $A^*$ is chosen optimally, it must be the case that expected output for the promoted workers is higher in the difficult job than in the easy job. If it were not, it would be better to raise $A^*$ until expected output were higher. Rather the point is that after promotion, the average promoted worker is not as able in the difficult job as he was in the easy job, that is,

$$\mathrm{E}(A + \epsilon_2 | \text{promoted}) < \mathrm{E}(A + \epsilon_1 | \text{promoted}).$$

Also true is that within any job, those left behind and not promoted have lower average ability than those of their cohorts entering into that job. If there were a series of promotion rounds, then at every level, those who were not promoted would have a job-specific component that is negative. This can be seen simply by examining the first round, which can be thought of as a "promotion" from being out of the firm to being hired as a worker. (Individuals must exceed some standard in order to be hired.) Since it has already been shown that

$$\mathrm{E}(\epsilon_1 | A + \epsilon_1 > A^*) > 0$$

and since $\mathrm{E}(\epsilon_1) = 0$, it must be true that

$$\mathrm{E}(\epsilon_1 | A + \epsilon_1 \leq A^*) < 0.$$

They appear "incompetent" because, within any given job, the actual ability of those who are not promoted out of the job is lower than the unconditional expectation of ability for that job. Those who are left behind and become the long-termers are worse than those who come into the job. They are incompetent relative to the entry pool because the competent workers are promoted out of the job. In a tournament with enough steps, each competent worker would continue to be promoted until he too was incompetent, that is, until $\mathrm{E}(\epsilon_t) < 0$ for those whose highest job attained is job $t$. This is the Peter principle: Workers are promoted to their level of incompetence. Those who are "competent" are promoted again.[10]

[10] One difference between this interpretation of the Peter principle and the one used in the rest of this paper is that the output of those not promoted does not rise under the job-specific interpretation. Since $\epsilon_1$ is a job-specific effect and not a transitory component, those who are not promoted have ability $A + \epsilon_1$ in both periods.

### D. Mistake or Optimal Adjustment?

The original book was entitled *The Peter Principle: Why Things Always Go Wrong*. The implications of the view that promotion decisions are biased are different from those of this model. The view in this model is that if $A^*$ is chosen optimally, most of the time firms will not want to undo their decision. Even though ability is below that predicted by a naive use of the first-period estimate, the promoted worker's ability is still above $x$, at least on average. If it were not, then the choice of $A^*$ would have been suboptimal. Ability falls but does not fall below $x$ in most cases, so the firm does not want to reverse its decision. The behavioral view that "things go wrong" is different. If the firm really made a mistake, then it would want to demote or fire workers in most cases, which leads to the inevitable question, "Why are demotions so rare?" This analysis provides an answer. The promotion rule is chosen optimally so that ability is not as high as it was before promotion, but it is still high enough to justify the promotion.

## II. Strategic Behavior by Workers

So far, worker effort has been assumed to be given. In this section, I relax the assumption that effort is given in order to determine how workers may game the system to alter their promotion possibilities. Whether workers overproduce during the probationary period depends crucially on the compensation scheme.

### A. Efficient Effort with Worker Job Choice

The first result is that if workers are paid piece rates and allowed to choose their own jobs, all is efficient, even if workers know their ability and firms do not. In order to examine incentives, it is necessary to define more terms: let $\mu_e$ be the amount of effort that an individual chooses if he is in the easy job and $\mu_d$ be the amount of effort he chooses if he is in the difficult job. Let $C(\mu)$ be the cost incurred for any given level of effort $\mu$. Then if $A$ is known but $\epsilon$ is not, individuals for whom

$$\int_{-\infty}^{\infty} [\gamma + \delta(A + \epsilon + \mu_d)]g(\epsilon)d\epsilon - C(\mu_d) >$$

$$\int_{-\infty}^{\infty} [\alpha + \beta(A + \epsilon + \mu_e)]g(\epsilon)d\epsilon - C(\mu_e) \quad (3)$$

choose the difficult job. Those for whom the condition in (3) does not hold choose the easy job. Effort levels in (3) are merely the optimal levels, given the job chosen.

Because the expectation of $\epsilon$ is zero, (3) can be rewritten as

$$\alpha + \beta(A + \mu_e) - C(\mu_e) < \gamma + \delta(A + \mu_d) - C(\mu_d). \tag{4}$$

If the condition in (4) holds, a worker prefers the difficult job. If it does not, the easy job is selected.[11] This is the same as the efficiency condition, so workers choose jobs and effort efficiently under these conditions. There is no distortion in effort choice. The worker internalizes everything. This is a simple problem of occupational choice with effort. By contrast, were the firm to choose the job for the worker, then effort would be distorted, although the surprising result is that the distortion is as likely to take the form of underwork as it is of overwork. This is shown in the Appendix.

### *B. Tournaments*

The usual intuition that most have about promotions inducing atypically high effort in period 1 comes from a tournament-like payment structure.[12] When period 2 wages depend on the job rather than on the output in the job, all workers put forth more effort than they would in the absence of period 2 promotion concerns.

The intuition holds whether the tournament is against another player or against a standard. For the purposes here, there is little difference between competing against another player and competing against a standard. In the Lazear-Rosen tournament structure, any level of effort can be implemented for any standard simply by choosing the wage spread appropriately. In the case of a standard, wages in period 2 are fixed in advance and depend only on promotion. Even if workers receive no wage prior to promotion, workers who are ignorant of $A$ put forth effort in order to maximize

$$\max_{\mu_1} \int \{W_d \Pr(A + \mu_1 + \epsilon_1 > A^*) + W_e[1 - \Pr(A + \mu_1 + \epsilon_1 > A^*)] - C(\mu_1)\} g(A)dA,$$

where $W_d$ is the difficult job's wage and $W_e$ is the easy job's wage.

[11] In a competitive market with a rising supply price for workers (because they are distinguished by ability), firms earn zero profit. The marginal worker is the one for whom ability $A_0$ is low enough that

$$\alpha + \beta(A_0 + \mu) - C(\mu) = 0.$$

[12] Here again, Rosen is instrumental. The first paper on the subject is Lazear and Rosen (1981).

The first-order condition is

$$\int \left[(W_d - W_e)\frac{\partial \Pr (A + \mu_1 + \epsilon_1 > A^*)}{\partial \mu_1} - C'(\mu_1)\right] dG = 0$$

or

$$(W_d - W_e) g(A^* - \mu_1 - A) = C'(\mu_1).$$

The firm can obtain any level of effort, $\mu_1$, simply by setting the spread between the difficult job wage and easy job wage appropriately. Then it is necessary only to set the expected wage sufficiently high to attract the marginal worker.

What is clear, however, is that effort in period 1 exceeds that in period 2. The tournament structure induces individuals to work at some positive level in period 1 but to reduce effort in period 2. In this stylized model, since there is no contingent reward in period 2, effort falls to zero. But the general point is that the tournament against a standard creates incentives to perform better in the prepromotion period than in the postpromotion period.

Firms understand that their compensation schemes induce strategic behavior by workers and set $A^*$ accordingly. Although this may mitigate the effects of the behavior, it in no way changes the results of this section. Since all derivations hold for any given $A^*$, they hold for the $A^*$ chosen to take these effects into account.

As is the case of the tournament against a standard, workers put forth more effort before the promotion decision than after the promotion decision in a tournament against another player. This follows directly from Lazear and Rosen (1981), where effort during the contest period exceeds effort after the contest period. Worker effort during the contest period is monotonically increasing in the spread between the winner's wage and the loser's wage. After the contest has been decided, effort falls off.

In both the tournament story and the regression to the mean story, worker output declines after the promotion decision. In the tournament context, the reason is that effort declines. In the regression to the mean version, output declines because of the statistical proposition that ensures that winners do worse after promotion. There is a difference, however. In tournaments, even losers reduce effort after the promotion has been decided, so expected output for all workers falls over time. In the statistical version, winners' output falls and losers' output rises above their prepromotion levels on average.

To amplify this point, just as those who are promoted have higher-than-average prepromotion transitory error, $\epsilon_1$, so do those who fail to be promoted have lower-than-expected transitory components. Other things equal, this implies that those who do not get a promotion should do better after being turned down than they did before. Thus faculty who are denied tenure and move to other schools should do better, on average, at those other schools than they did when they were assistant professors at the first institution.

Observing this effect may be difficult for a number of reasons. For example, a worker's output might depend on the individuals with whom he works. In an up-or-out context (Kahn and Huberman 1988), those who fail to be promoted may find that the complementary factors in the new job are not as productivity-enhancing as those in the first job. Furthermore, motivation is an issue. To the extent that an individual believes that he is in the running for promotion, tournament effects are present, inducing effort. After the promotion has been denied, the incentives vanish, reducing effort and output.

## III. Other Examples of the Principle

The regression to the mean phenomenon that is observed as the Peter principle in the labor market has other manifestations. For example, it is often observed that sequels are rarely as good as the original movie on which the sequel is based.[13] If each movie is thought of as having a theme-constant component, $A$, and a transitory component, $\epsilon_t$ (e.g., actors, specific story, or direction), associated with each particular film, then the same analysis holds. In order for a sequel to be made, the value of the original film, $A + \epsilon_t$, must be estimated to be greater than $A^*$. But given that the value exceeds the threshold level, $A^*$, the expectation of the value of the sequel will be less than that of the original simply because

$$\mathrm{E}(\epsilon_1|\text{sequel is made}) > 0,$$

but

$$\mathrm{E}(\epsilon_2|\text{sequel is made}) = 0.$$

As a result, an original film must be sufficiently good to generate a sequel because studios, knowing that the second film is likely to be inferior to the first, adjust upward their cutoff levels.

It is straightforward to test this proposition. Among other things, it

[13] Data from Teichner and Luehrman (1992) establish clearly that revenues are lower and costs higher on sequels than they are on the original on which they are based. See, e.g., their exhibits 4 and 5.

implies that measures of film quality such as academy awards or ticket sales should be higher on the original film than on the sequels.[14]

The fact that sequels do worse than originals is not evidence that the studio made a mistake. Once again, if $A^*$ is chosen optimally, then the average sequel is profitable, even if not as profitable as the original. A decline in profitability does not imply that something went wrong in decision making.

Similarly, the first meal in a good restaurant is often the best, followed by less satisfying repeat visits. Just as above, think of $A$ as the restaurant component of the first meal (recipes, management) and $\epsilon_1$ as the transitory component associated with the meal itself (that night's chef, specific ingredients, dinner companion). A second visit to the restaurant is made only if $A + \epsilon_1 > A^*$. Once again, the expected value of the second meal lies below that of the first, conditional on the decision to make a second visit to the restaurant. The larger the transitory component, the larger the discrepancy between the first and second meals and the higher the standard set to merit a second visit. As before, the fact that the second visit is not as good as the first does not imply that a mistake was made. The second meal, on average, is good enough to justify a repeat visit to the restaurant, even if it is not as good as the first. The point can also be used to explain why favorite restaurants go out of fashion. A restaurant becomes a favorite in part because of the permanent component (e.g., good recipes and an insightful owner) and in part because of potentially transitory components (e.g., the current chef and the service of the staff). A favored restaurant can be thought of as one that has gotten a draw of $A + \epsilon_1 > A^*$. It is favored precisely because the value of the output exceeds some standard. Over time, $\epsilon_1$ is replaced by transitory effects, the expectation of which is zero. The quality falls and the restaurant goes out of fashion.

The same logic can be applied to the "*Sports Illustrated* effect." It is claimed that it is a curse to be on the cover of *Sports Illustrated* because athletes' performance falls thereafter. Again this reflects regression to the mean, but it does not imply that *Sports Illustrated* chose the wrong athlete for the cover.

In the political arena, reelections of elected officials occur when constituents view performance as having exceeded some standard. The postelection performance should be worse without learning and seniority effects. If the learning effect is not too strong, second-term presidents will do less well than they did during their first term. But this does not imply that voters made a mistake. Even when voters know that there is

[14] A countervailing effect is the notoriety that is created by the first film, which makes it easier to sell tickets on the sequel than on the original. Even if consumers understand that the sequel is worse than the original, more tickets might be sold on the sequel if, say, the actors and director are not well known before the first film is made.

regression, voting for the incumbent is rational if he exceeds a sufficiently high standard.

Finally, there is a close relation of this analysis to that of the winner's curse (Wilson 1969).[15] If a bidder submits a bid based on a naive estimate, the winner will lose money because, on average, the winner has obtained a higher-than-average error. As a result, he shades his bid, so in equilibrium he adopts the rent-maximizing strategy. The same is true here. By adjusting the cutoff criterion, $A^*$, the decision maker maximizes rent.

## IV. Length of Probationary Period and Relative Importance of the Transitory Component

The longer a firm waits to make a promotion decision, the better the information. One would expect that transitory components that bias a decision could be reduced or eliminated if the firm waited long enough to make a promotion decision. The cost of waiting, however, is that workers are in the wrong job for more of their lifetimes. For example, suppose that it were possible to get a perfect reading on $A$ by waiting until the date of retirement. The information would have no value because the worker would have spent his entire working career in the easy job, even if he were better suited to the difficult job. The trade-off is modeled. The conclusion is that as the variance of $\epsilon_1$ rises, it becomes more valuable to wait on a promotion decision.

To see this, let us add one period to the previous model (without effort). Now, $\epsilon_1$, $\epsilon_2$, and $\epsilon_3$ refer to the transitory component in periods 1, 2, and 3, and assume that they are distributed independently and identically and, to reduce notation, that $E(A) = 0$. Suppose that by waiting two periods, an employer can obtain a perfect reading of $A$. Under those circumstances, the optimum is simply to promote those for whom $A > x$. The cost is that when the firm delays its promotion decision to the end of period 2, all workers are in the easy job during period 2 even though it might be better to place some in the difficult

[15] Winner's curse usually relates to a reading relative to others' readings rather than the time dimension of taking multiple readings, sometimes in different settings. The transitory vs. permanent component is central to the discussion of this paper, but not the theme of most of the winner's curse literature. Actually, loser's curse is as important to the assignment problem as winner's curse. In the job context, the goal is to assign a worker to the right job. Workers who do not satisfy the promotion criterion are, on average, undervalued just as those who are promoted are overvalued. The optimal selection of the cutoff point trades off the two kinds of errors.

job in period 2. Expected output over the lifetime is then

$$\begin{aligned}\text{expected output if wait} \\ &= 2[\alpha + \beta\mathrm{E}(A + \epsilon_1)] + \int_x^\infty (\gamma + \delta A)d + \int_{-\infty}^x (\alpha + \beta A)dFF \\ &= 2\alpha + \int_x^\infty (\gamma + \delta A)dF + \int_{-\infty}^x (\alpha + \beta A)dF. \qquad (5)\end{aligned}$$

The alternative is to make a decision after one period, using imperfect information and recognizing that sorting will be imperfect. To make things simple, assume that a firm that makes a promotion decision at the end of period 1 cannot reevaluate at the end of period 2. The gain is that workers are sorted early so that very able people can be put in the difficult job more quickly. The cost is that more errors are made in assigning workers to jobs. Then expected output over the three periods is

$$\begin{aligned}\text{expected output early} \\ &= [\alpha + \beta\mathrm{E}(A + \epsilon_1)] + 2\int_{-\infty}^\infty \int_{A^*-\epsilon}^\infty (\gamma + \delta A)dFdG \\ &\quad + 2\int_{-\infty}^\infty \int_{-\infty}^{A^*-\epsilon} (\alpha + \beta A)dFdG \\ &= \alpha + 2\int_{-\infty}^\infty \int_{A^*-\epsilon}^\infty (\gamma + \delta A)dFd + 2\int_{-\infty}^\infty \int_{-\infty}^{A^*-\epsilon} (\alpha + \beta A)dFdG. \qquad (6)\end{aligned}$$

In an extreme case, it is clear that it pays to decide early. If the distribution of $\epsilon$ is degenerate so that there is no error, then (6) becomes

$$\text{expected output early} = \alpha + 2\int_x^\infty (\gamma + \delta A)dFdG + 2\int_{-\infty}^x (\alpha + \beta A)dF. \qquad (7)$$

The right-hand side of the expression in (7) must exceed the right-hand side of (5) because

$$\gamma + \delta A > \alpha + \beta A \quad \text{for } A > x$$

since that is how $x$ is defined. Thus, when the variance in $\epsilon$ shrinks to zero, it always pays to promote early.

The example used earlier shows that it sometimes pays to defer the promotion decision until the end of period 2. As before, let $\alpha = 1$,

$\beta = 0.5$, $\gamma = 0$, and $\delta = 1$, where the distributions of $A$ and $\epsilon_t$ are normal with variance equal to one. As shown earlier, the optimal cut point is $A^* = 4.01$. Then the right-hand side of (5) equals 3.004. The right-hand side of (6) is 3.000. Thus deferring the promotion decision until the second period pays when the variance in $\epsilon$ is one. Other numerical examples show that the advantage of deferring the promotion decision becomes larger for higher variances in $\epsilon$.

The general point is that when the distribution of $\epsilon$ is sufficiently tight, it pays to make the promotion decision early. When it is sufficiently diffuse, it pays to make the promotion decision later. Later promotion decisions are more accurate but result in workers' spending a longer proportion of their work life in the wrong job.

## V. Conclusion

Workers who are promoted have been observed to have exceeded some standard. Part of the observation is based on lasting ability, but part is based on transitory components that may reflect measurement difficulties, short-term luck, or skills that are job-specific. As a result, there is regression to the mean, creating a "Peter principle." Workers who are promoted do not appear to be as able as they were before the promotion.

Firms take this phenomenon into account in setting up their promotion rule. Under general conditions, when fewer than 50 percent of the workers are better suited to the high-level job, the firm adjusts the promotion standard upward to compensate for the regression to the mean. The amount of the adjustment depends on the tightness of the error distribution. When the prepromotion error has high dispersion, promotion standards are inflated by more than they are when the error dispersion is low.

The statistical argument has been contrasted with incentive arguments. Whether workers overproduce because they are gaming the system depends on the payment structure. If, for example, output were observable so that workers could be paid on the basis of output both before and after the promotion decision, then it would be optimal to allow workers to make their own job choice. Under these circumstances, there is no distortion in effort; all is efficient. When a tournament structure is chosen because of inability to observe output, workers produce more before promotion than they do after promotion. Although tournaments result in declining output after promotion, the implications of tournaments for losers and winners are different from those of the statistical argument. In particular, in tournaments, the output of both losers and winners falls after promotion. The statistical argument implies that losers' output rises and winners' output falls after promotion.

The Peter principle can be interpreted to mean that workers are not as able as perceived before promotion or that they were better in their prior job relative to their peers than they are in their current one. In a multilevel firm, the typical worker who remains at a given level is "incompetent" in that he is not as good as the average worker coming into the job, nor is he as good as he was in his previous assignment relative to the comparison set.

One way to offset the Peter principle is to wait for a longer time before making a promotion decision. The advantage is that the job assignment is better than it would have been had the decision been made earlier. The disadvantage is that able workers remain in the wrong job for a longer period of time.

The logic of the Peter principle applies in other contexts as well. The regression to the mean phenomenon implies that movie sequels are lower-quality than the original films on which they are based and that excellent restaurant meals are followed by ones that are closer to mediocre.

## Appendix

### Effort Is Distorted When Firms Assign Jobs but Workers Know Ability

Assume asymmetric information in which workers know their abilities, $A$, but firms do not. Define $\mu_1$ as effort in period 1, $\mu_{2e}$ as effort in period 2 if the worker is not promoted, and $\mu_{2d}$ as effort in period 2 if the worker is promoted. Note that effort in period 1 is determined before the promotion decision is made, so period 1 effort is independent of promotion. The cost of effort is given by $C(\mu)$. For simplicity, $C(\mu)$ is assumed to be independent of ability and the same across periods.

The worker is paid a piece rate, so in period 2, a worker who has not been promoted chooses effort $\mu_{2e}$ so as to solve

$$\max_{\mu_{2e}} \alpha + \beta \mathrm{E}(A + \mu_{2e} + \epsilon_2) - C(\mu_{2e})$$

or

$$\max_{\mu_{2e}} \alpha + \beta(A + \mu_{2e}) - C(\mu_{2e}).$$

The first-order condition is

$$C'(\mu_{2e}) = \beta. \tag{A1}$$

An analogous problem can be solved for those who are promoted. Their problem is

$$\max_{\mu_{2d}} \gamma + \delta(A_2 + \mu_{2d}) - C(\mu_{2d}),$$

which has first-order condition

$$C'(\mu_{2d}) = \delta. \tag{A2}$$

Equations (A1) and (A2) define $\mu_{2e}$ and $\mu_{2d}$. Promoted workers put forth more effort in period 2 because the marginal return to effort is higher in the difficult job than in the easy job, that is, $\delta > \beta$. Given this, the worker solves a two-period problem in period 1, knowing that he will choose $\mu_{2d}$ and $\mu_{2e}$ depending on whether or not he is promoted.

The worker who knows his own ability has a first-period problem given by[16]

$$\max_{\mu_1} \alpha + \beta \mathrm{E}(\mu_1 + A + \epsilon_1) - C(\mu_1)$$
$$+ \operatorname{Prob}(A + \mu_1 + \epsilon_1 > A^*)\mathrm{E}[\gamma + \delta(A + \mu_{2d} + \epsilon_2) - C(\mu_{2d})]$$
$$+ \operatorname{Prob}(A + \epsilon_1 + \mu_1 \leq A^*)\mathrm{E}[\alpha + \beta(A + \mu_{2e} + \epsilon_2) - C(\mu_{2e})]$$

or

$$\max_{\mu_1} \alpha + \beta(\mu_1 + A) - C(\mu_1) + [1 - G(A^* - \mu_1 - A)][\gamma + \delta(A + \mu_{2d})$$
$$- C(\mu_{2d})] + G(A^* - \mu_1 - A)[\alpha + \beta(A + \mu_{2e}) - C(\mu_{2e})]. \tag{A3}$$

The first-order condition is

$$\beta - C'(\mu_1) + g(A^* - \mu_1 - A)\{[\gamma + \delta(A + \mu_{2d}) - C(\mu_{2d})]$$
$$- [\alpha + \beta(A + \mu_{2e}) - C(\mu_{2e})]\} = 0. \tag{A4}$$

Efficient effort is supplied when workers set $C'(\mu_1) = \beta$. According to the first-order condition in (A4), this occurs only when the last term on the left-hand side is equal to zero. In general, it will not be zero. In fact, the last term is positive, implying overinvestment, when

$$\gamma + \delta(A + \mu_{2d}) - C(\mu_{2d}) > \alpha + \beta(A + \mu_{2e}) - C(\mu_{2e}).$$

Sufficiently high-ability workers prefer job 1 because they earn more in job 1. As a result, they overwork in period 1 to enhance the probability that they will be promoted. Because the firm cannot distinguish effort from ability, workers who want to be promoted have an incentive to work too hard in order to fool the firm into believing that their ability levels are higher than they actually are.

Less intuitive, the converse is also true. Low-ability workers, that is, those for whom $A$ is sufficiently low so that

$$\gamma + \delta(A + \mu_{2d}) - C(\mu_{2d}) < \alpha + \beta(A + \mu_{2e}) - C(\mu_{2e}),$$

underwork.[17] These workers underachieve because they do not want to take the chance of being promoted. From their point of view, a promotion is bad because they are likely to earn less in the difficult job than in the easy job.

The workers who are most likely to distort their effort in period 1 are those for whom $g(A^* - \mu_1 - A)$ is high (see eq. [A4]) and for whom

$$|[\gamma + \delta(A + \mu_{2d}) - C(\mu_{2d})] - [\alpha + \beta(A + \mu_{2e}) - C(\mu_{2e})]|$$

[16] The discount rate is assumed to be zero.

[17] Since $\mu_2$ is independent of $A$, there is always an $A$ sufficiently low to make this condition hold.

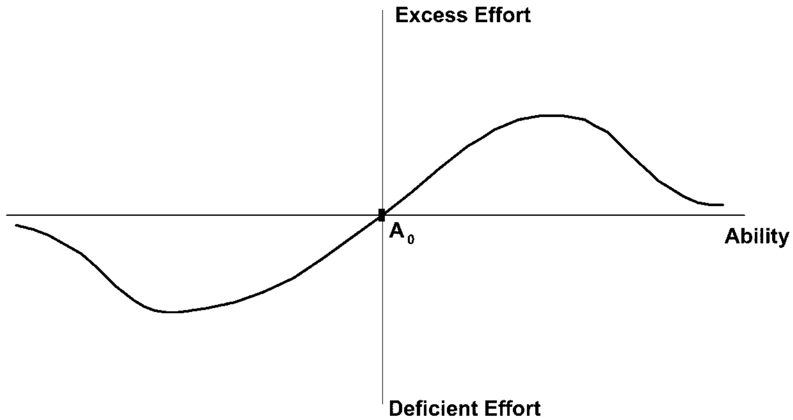


Fig. A1

is high. Under standard assumptions about the distribution of $\epsilon$, in particular that

$$\lim_{\epsilon \to -\infty, \infty} g(\epsilon) = 0,$$

very high- and very low-ability workers choose the efficient level of effort in period 1. They have little to fear in terms of incorrect promotion decisions. The extremely able are almost certain to be promoted, so that extra effort has very little effect on the probability of promotion. Conversely, the totally inept are almost certain to avoid promotion, so that reducing effort has almost no effect on lowering the probability of promotion.

Also true is that those whose underlying ability is very near the efficient job switch point ($x$ in fig. 1) do not distort effort much. Even if they are misclassified, they have little to lose. Define $A_0$ such that

$$\gamma + \delta(A_0 + \mu_{2d}) - C(\mu_{2d}) = \alpha + \beta(A_0 + \mu_{2e}) - C(\mu_{2e}).$$

Then excess effort is zero at both extremes and also at $A_0$, which is likely to be close to $x$. The pattern of distortion is shown in figure A1. Those at $A_0$ do not distort at all. Those at the ability extremes do not distort. Those with ability less than the switch point underwork and those with ability more than the switch point overwork.

## References

Anderson, Rolph E., Alan J. Dubinsky, and Rajiv Mehta. 1999. "Sales Managers: Marketing's Best Example of the Peter Principle?" *Bus. Horizons* 42 (January/February): 19–26.

Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. "The Internal Economics of the Firm: Evidence from Personnel Data." *Q.J.E.* 109 (November): 881–919.

Fairburn, James A., and James M. Malcomson. 2001. "Performance, Promotion, and the Peter Principle." *Rev. Econ. Studies* 68 (January): 45–66.

Faria, Joao Ricardo. 2000. "An Economic Analysis of the Peter and Dilbert Principles." Manuscript. Sydney: Univ. Technology.

Gibbons, Robert, and Michael Waldman. 1999*a*. "Careers in Organizations: Theory and Evidence." In *Handbook of Labor Economics*, vol. 3B, edited by Orley C. Ashenfelter and David Card. Amsterdam: North-Holland.

———. 1999*b*. "A Theory of Wage and Promotion Dynamics inside Firms." *Q.J.E.* 114 (November): 1321–58.

Gibbs, Michael, and Wallace Hendricks. 2001. "Are Formal Salary Systems a Veil?" Manuscript. Chicago: Univ. Chicago.

Jovanovic, Boyan. 1979*a*. "Firm-Specific Capital and Turnover." *J.P.E.* 87 (December): 1246–60.

———. 1979*b*. "Job Matching and the Theory of Turnover." *J.P.E.* 87, no. 5, pt. 1 (October): 972–90.

Kahn, Charles, and Gur Huberman. 1988. "Two-Sided Uncertainty and 'Up-or-Out' Contracts." *J. Labor Econ.* 6 (October): 423–44.

Lazear, Edward P. 1986. "Raids and Offer Matching." In *Research in Labor Economics*, vol. 8, pt. A, edited by Ronald G. Ehrenberg. Greenwich, Conn.: JAI.

———. 1992. "The Job as a Concept." In *Performance Measurement, Evaluation, and Incentives*, edited by William J. Bruns, Jr. Boston: Harvard Bus. School Press.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *J.P.E.* 89 (October): 841–64.

Medoff, James L., and Katharine G. Abraham. 1980. "Experience, Performance, and Earnings." *Q.J.E.* 95 (December): 703–36.

Peter, Lawrence J., and Raymond Hull. 1969. *The Peter Principle: Why Things Always Go Wrong.* New York: Morrow.

Prendergast, Canice. 1999. "The Provision of Incentives in Firms." *J. Econ. Literature* 37 (March): 7–63.

Rosen, Sherwin. 1986. "Prizes and Incentives in Elimination Tournaments." *A.E.R.* 76 (September): 701–15.

Teichner, William A., and Timothy A. Luehrman. 1992. "Arundel Partners: The Sequel Project." Harvard Business School Case no. 9-292-140. Boston: Harvard Bus. School Pub.

Wilson, Robert B. 1969. "Competitive Bidding with Disparate Information." *Management Sci.* 15 (March): 446–48. Reprinted in *The Economics of Information*, vol. 1, edited by David K. Levine and Steven A. Lippman. Aldershot, U.K.: Elgar, 1995. Also reprinted in *The Economic Theory of Auctions*, vol. 2, edited by Paul Klemperer. Cheltenham, U.K.: Elgar, 2000.